

EvoTADASHI: Genetic Programming for C Code Optimization

João Eduardo Batista, Emil Vatai, Aleksandr Drozd, Mohamed Wahib

RIKEN Center for Computational Science, Kobe, Japan



Introduction

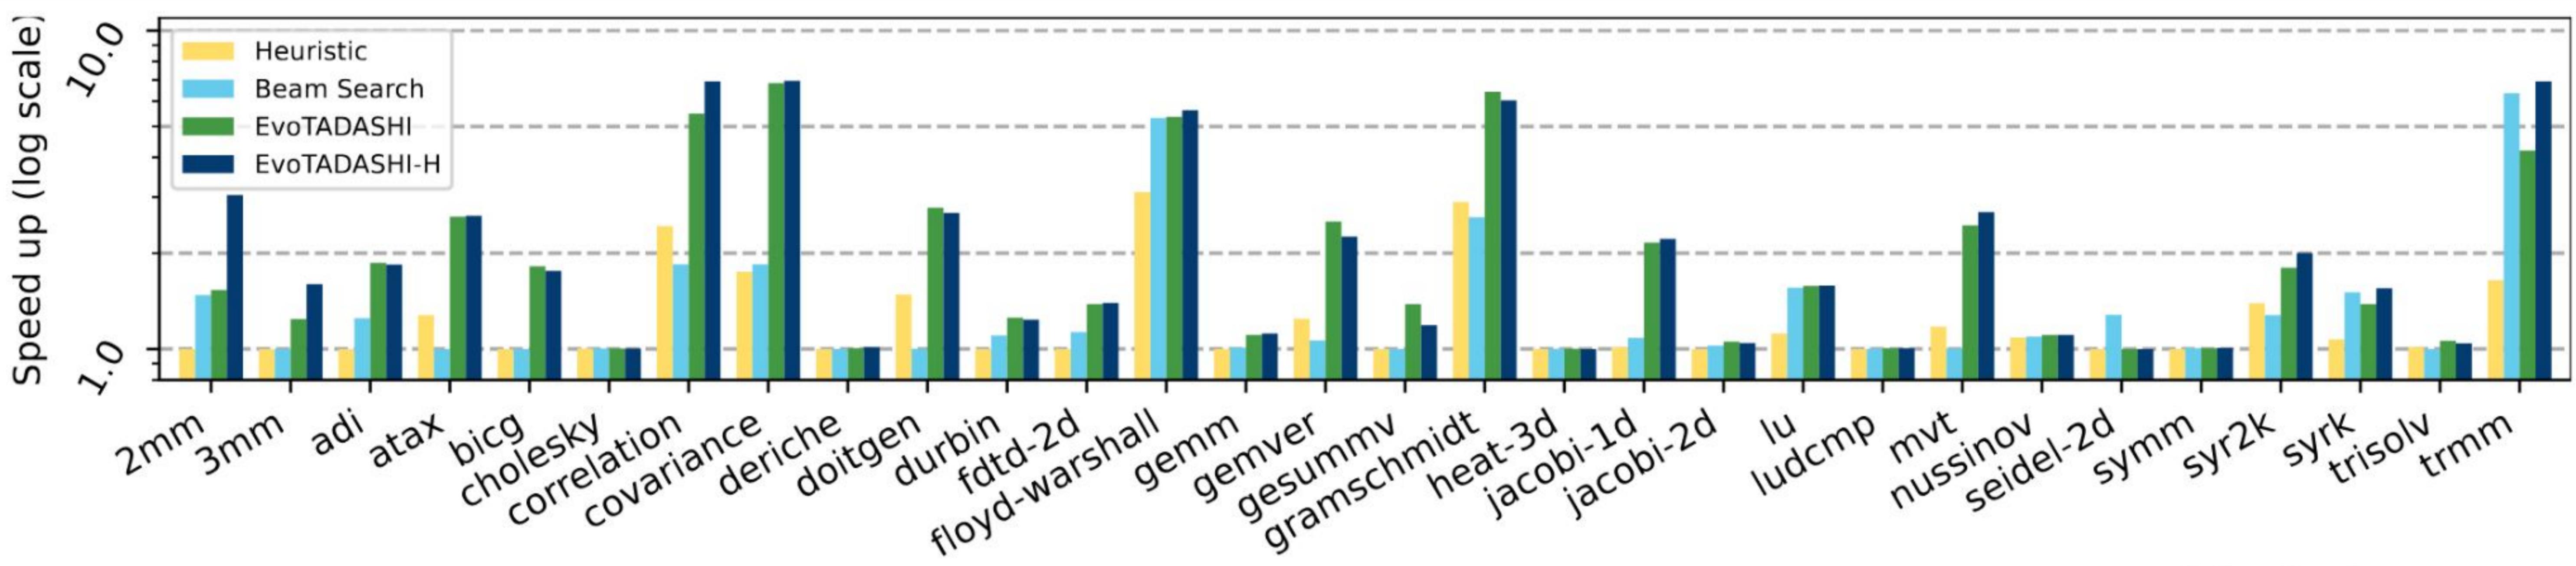
When optimizing code, the code must be adapted to the target hardware every time a new system is bought. This is typically done manually and is both time-consuming and error-prone. Machine learning has automated some aspects of this process, but typical black-box methods lack guarantees of correctness or transformation legality.

TADASHI

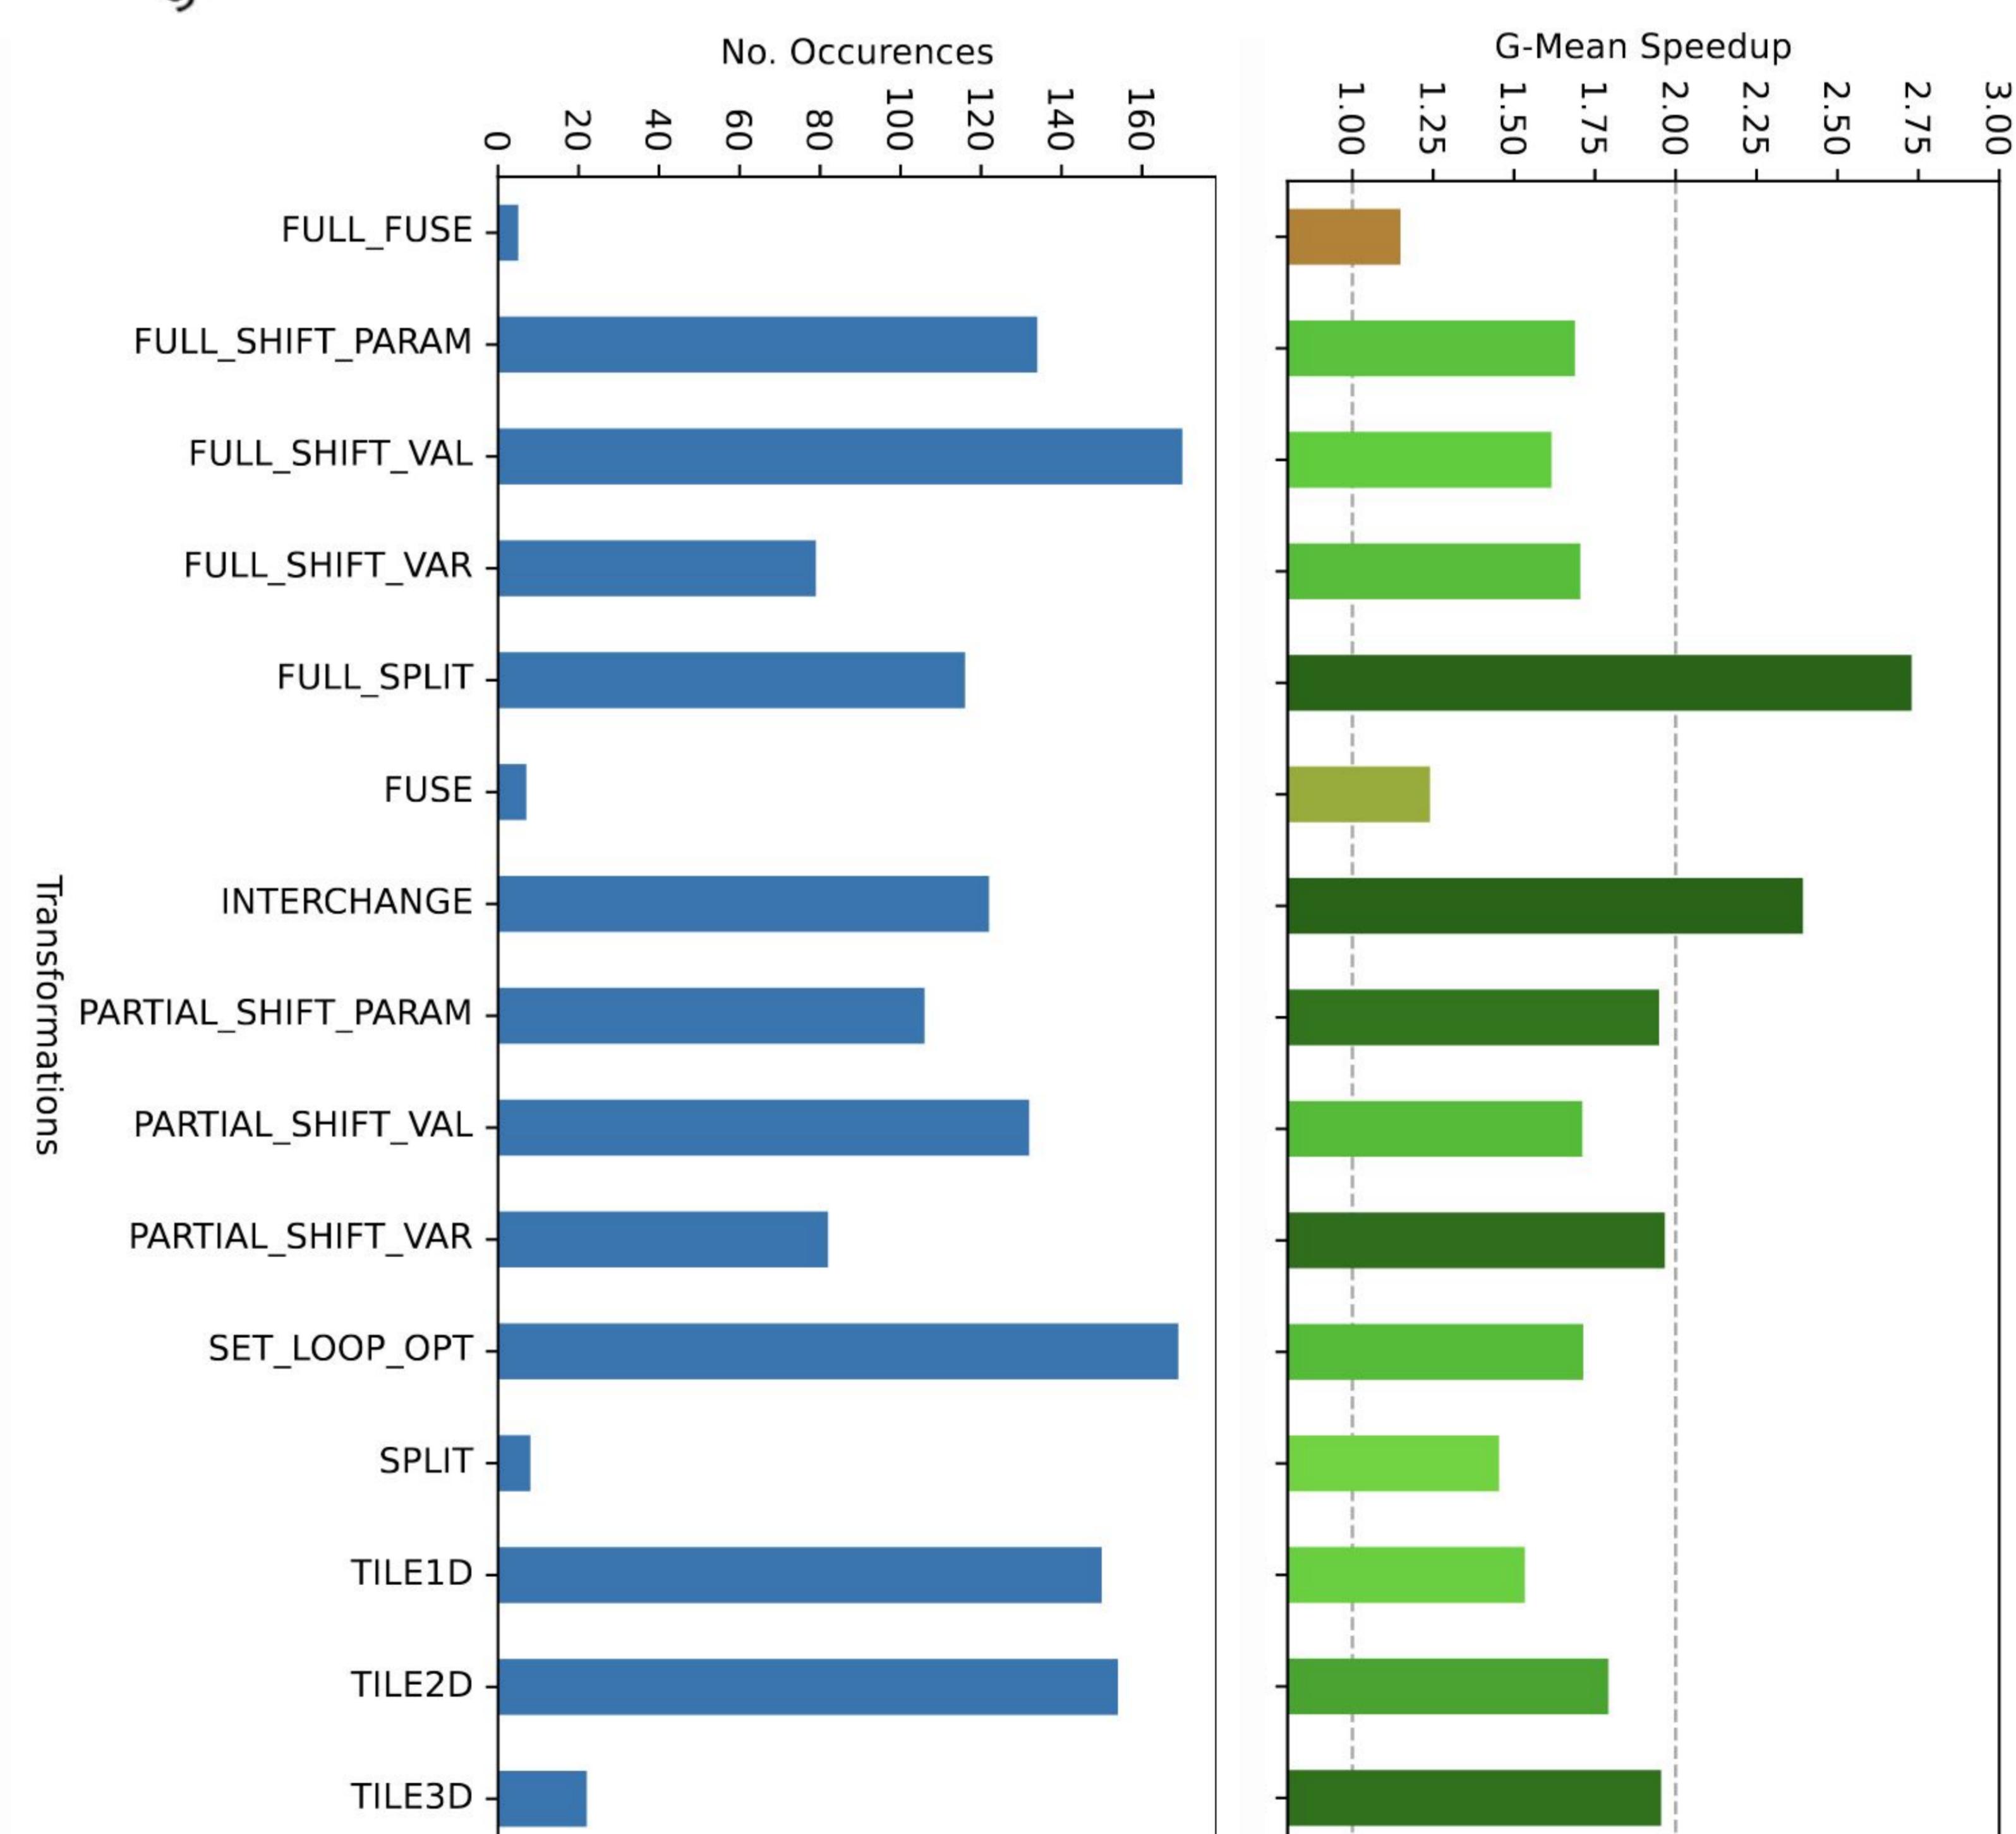
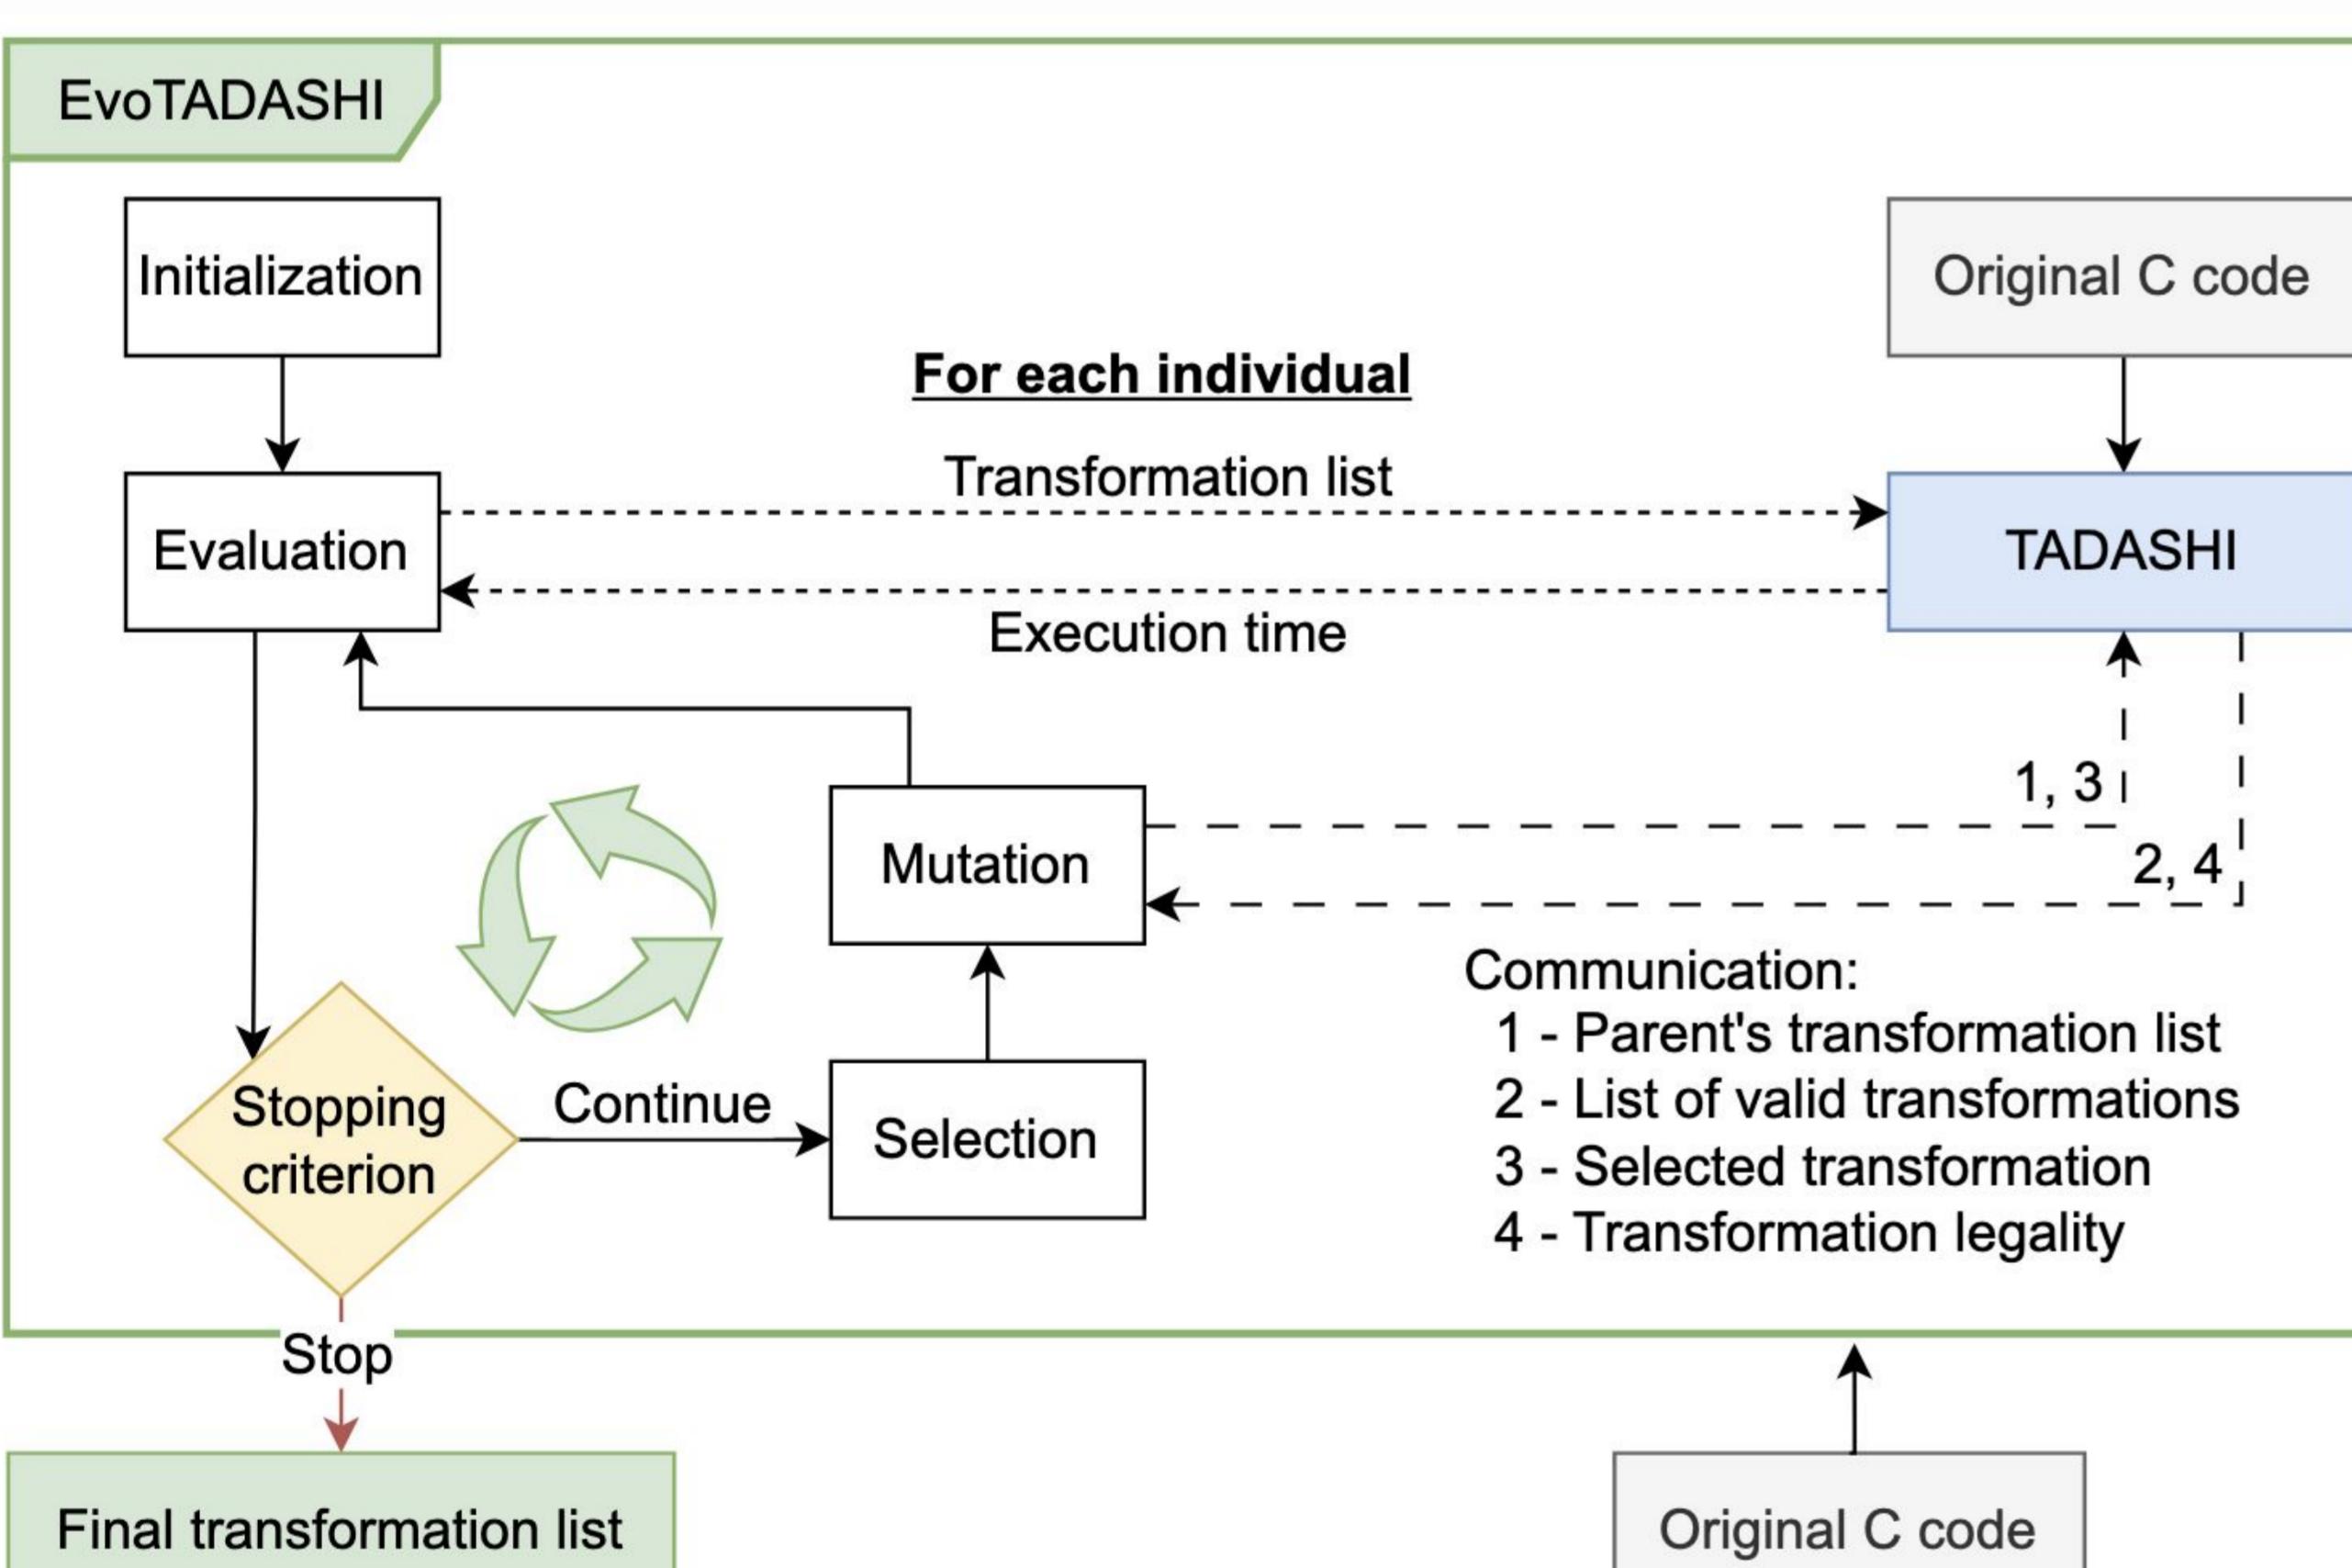
TADASHI is a Python library that bridges the gap between compiler technologies and machine learning. It enables machine learning experts with little to no compiler experience to transform and optimize their code while generating a list of transformations that clearly explain each modification step by step, exemplified on the right.

```
solution = [
    [45, FULL_SHIFT_VAR, 0, 23],
    [28, PARTIAL_SHIFT_PARAM, 1, 1, -14],
    [41, FULL_SPLIT],
    [3, FULL_SHIFT_PARAM, 0, -39],
    [45, INTERCHANGE]
]
```

Results



EvoTADASHI



Methodology

We introduce EvoTADASHI, a machine learning algorithm based on genetic programming that optimizes C code using the TADASHI library. We also compare the results obtained by EvoTADASHI with those obtained using the heuristic proposed in the original TADASHI paper; with our implementation of beam search, and with EvoTADASHI-H, an approach that bootstraps EvoTADASHI using the heuristic to initialize training. These methods are tested using PolyBench, a set of 30 C code benchmarks. The benchmarks were compiled with the extra-large dataset and -O3 gcc (v11.5.0) flags on nodes with an AMD EPYC 9684X (Genoa-X) 96-Core 2.55 GHz processors with a 3.7 GHz max boost, a large 1152 MB of L3 cache, and 768GB of RAM.

Conclusions and Future Work

Our results show that base EvoTADASHI matches beam search's performance (1.86x speedup) with half the computational cost. The simpler heuristic method (1.18x) still improved 11 benchmarks in a single step. When used to initialize EvoTADASHI, performance increased further to 2.04x. Overall, EvoTADASHI-H delivered speedups in 27 of 30 benchmarks.